AI ATAC #2 - FAQ

Q1: I am a security architect and I would like to see the list of participants and the results of this competition. Would that be possible?

A1: No, we cannot share the list of participants and the results of the competition. We will share who won the competition, but will not share their results.

Q2: We may work in conjunction with a VAR. Can we have them submit on our behalf? In that case, the primary POC would work for their company.

A2: Yes, per the eligibility section on the Challenge website:

"Participants must either own the intellectual property (IP) in the solution or provide documentation demonstrating exclusive arrangements and/or rights with the IP owner".

Q3: We are planning on participating in the challenge and will likely also include an SI partner and technology partner, both of which are TBD. Will you need an update before we submit on the 29th of May?

A3: Yes, you must notify us before submitting your solution. Please note the challenge submission date has changed to June 30th, 2020.

Q4: In the eligibility criteria, it states you must be incorporated in the US or have a place of business in the US. We are currently registered in SAM (https://sam.gov/SAM/) and DLA (https://www.dla.mil/). Would we be eligible to participate in this challenge or in future challenges?

A4: Per the rules, the participant must be:

1) incorporated in the U.S., and
2) maintain a primary place of business in the U.S.

So long as both of these is satisfied, this eligibility aspect is satisfied.

Q5: One requirement states that technologies will be provided a passive tap for the enterprise network. If traffic is being fed from a passive tap, how are vendor technologies expected to decrypt SSL traffic? Please clarify if different wording should be used here or if there are additional architectural concepts to consider.

A5: The tool will be tested on both encrypted and unencrypted traffic and SSL decryption capabilities will not be permitted and are out of scope for the challenge.

Q6: What is the proposed % of encrypted vs unencrypted traffic?

A6: This information will not be shared.

Q7: Can multi-vendor solutions be submitted? For example, if a single vendor solution doesn't provide SSL decryption or load balancing natively, can another vendor be used to accomplish this task?

A7: SSL decryption capabilities will not be permitted and are out of scope for the challenge.

The Evaluation Team will make the best effort to accommodate load balancing technologies, if needed. Network level detection technologies can partner with load balancing capabilities, but this requires written consent by the third parties for the use of these complementary capabilities. The network detection technology is what is eligible for award and evaluation. The third party is not eligible for award in this case. These third parties could be used by multiple eligible network detection technologies in the AI ATAC 2 Challenge (with their consent for each).

Q8: If load balancing is required to ensure all packets are processed:

a. **Which load balancing solution will ORNL use?**
b. **Or, is the vendor expected to supply their own or 3rd party load balancing solution?**

A8: If load balancing is needed to support up to 10Gb/s bandwidth, the stated maximum traffic rate from the challenge guidelines, then the submission technology should provide all necessary components. See question #7 above about partnering with third-party load-balancing technologies.

Q9: What is the proposed throughput for East/West traffic?

A9: This information will not be shared.

Q10: Can we know what the traffic profile (network load wise) for east / west versus north south will be? 9GB internal / 1GB external?

A10: This information will not be shared.

Q11: Please elaborate on this statement: "Technologies that leverage learning online or from historical data will be provided three weeks of ambient network data to form their models." Does this mean that all vendors will be provided network data? If so, when? Will this data be sent offline before the competition, or is there a three-week period (of simulated time) of benign traffic for 'burn-in' purposes?

A11: Once submitted tools are in place—properly configured and locked for the evaluation in ORNL's Cybersecurity Operations Research Range (CORR)—approximately three weeks of ambient network data will be replayed to allow tools to "burn in", i.e., time to train to the evaluation environment's emulated network, if needed. This training is a pre-step to AI ATAC 2 evaluation, i.e., before adversarial campaign detection experiments begin.

Q12: Is there a rack space or power limitation?

A12: If special circumstances are expected, please include a description in the whitepaper. ORNL will work with accepted whitepaper submissions to make sure accommodations are met. If there is a specific rack space or power (or other) requirement that is exorbitant, e.g., more than 7U of rack space or more than 4 kW, please email these specific requirements ASAP to the AI ATAC group email address.

Q13: What is the physical connection of the passive tap by which a solution receives network traffic?

a. **A single connection? Copper/Fiber?**
b. **Multiple connections with east/west traffic in one connection and north/south in another?**
c. **Multiple connections, traffic could be asymmetric?**

A13: It will be a single connection for all traffic. Specific connection details will be worked out during the Technology Installation and Configuration phase. After whitepaper down-selection, eligible technologies will be configured and locked during the on-site vendor visit, or by the ORNL team if the vendor chooses not to visit and instead send the submission along with instructions for installation.

Q14: Can we assume that the traffic profile will contain standard application file types and in realistic proportions? For example, some third-party testing have sent enormous amounts of benign unknown files to disrupt or delay analysis, but is unrealistic in a production environment with real end-users.

A14: Adversarial campaigns will be based on real-world adversarial capabilities. Specific traffic and file distributions will not be shared.

Q15: One requirement states that technologies will be provided a passive tap for the enterprise network. Does that mean that the technologies will be tested in a passive, non-blocking mode?

A15: Yes, tested tools will be in a passive, non-blocking, only detecting and alerting mode.

Q16: Given the length of the challenge, does ORNL expect to manually update vendor ML engines or OS/Firmware at monthly/quarterly intervals or is the vendor submission "frozen in time" at initial deployment?

A16: For the submitted technologies that pass the whitepaper down-selection, the technologies "will be frozen in time" once they are properly configured in the testing environment. From the AI ATAC 2 call:

> "the evaluation team will attempt to install and configure each technology in the range environment, in order to assure communication with the range data interfaces and tune each technology's configuration. A maximum two-day on-site support by the submitting organization may be required to optimize the configuration. At the conclusion of the Installation and Configuration stage, the technology configuration will be locked"

No updates are possible after the technology is "locked. The technology will have ~3 weeks of ambient data to "burn in", i.e., time to train in the testing environment's network with network traffic similar to the evaluation environment.

Q17: What protocols are being tested? The challenge only says n/s + e/w encrypted and decrypted traffic. Is it the same traffic protocols as the previous test (HTTP protocol, email protocol, SMB, custom sockets, custom UDP)? If so, what email protocols specifically and which versions of SMB?

A17: A variety of protocols will be tested, and specifics will not be shared.

Q18: What is the rate at which files will be submitted and will various file types come in all at the same time or will there be a random spread of file types during the course of the testing?

A18: This will not be shared.

Q19: Once configured and updated, how long is our submission in offline mode before testing occurs? Is it the same offline time period for all vendors?

A19: See response to Question #16. For those submitted technologies that pass the whitepaper down-selection, the technologies "will be frozen in time" once they are properly configured in the testing environment. From the AI ATAC 2 call:

> "the evaluation team will attempt to install and configure each technology in the range environment, in order to assure communication with the range data interfaces and tune each technology's configuration. A maximum two-day on-site support by the submitting organization may be required to optimize the configuration. At the conclusion of the Installation and Configuration stage, the technology configuration will be locked".

No updates are possible after the technology is "locked". The technology will have ~3 weeks of ambient data to "burn in", i.e., time to train to the testing environment's network. This is the configuration that will be used for testing. The timeline for all participating technologies will be the same.

Q20: Do we ship physical Hardware for the submission deadline or wait to hear if we are down selected to participate?

A20: Yes, you must submit your white paper and your tool (hardware/software or some combination thereof) by June 30th, 2020.

Q21: Will the 10GB/s networks be fully saturated, and would it be 10Gb/s incoming & another 10Gb/s outgoing speed?

A21: Submissions should support up to 10GB/s total traffic (incoming and outgoing combined will be no more than 10GB/s). Said another way, north-south (internal-to/from-external) and east-west (internal-to/from-internal) traffic will be up to 10Gb/s. Since we will not disclose details that would allow candidate technologies to tune their solution directly to the test environment, we recommend providing both VMs and hardware if the candidate technology may need stronger performance, and the evaluation team can rely on the hardware if needed.

Q22: Will there be a formal invitation to ship HW if required where we can discuss configuration details?

A22: You must submit it per the instructions listed on the Challenge.gov website

Q23: What is the total amount of traffic Mb/s l Gb/s for Ingress / Egress traffic

A23: We will not disclose specific traffic numbers.

Q24: What is the total amount of traffic Mb/s l Gb/s for Internal / Internal traffic

A24: We will not disclose specific traffic numbers.

Q25: Will traffic monitoring be split between east / west l north / south to individual network sensors or a single monitoring interface?

A25: All traffic will be mirrored with a SPAN and a copy of the traffic sent to the tool. The N-S and E-W traffic will be mirrored to a single monitoring interface.

Q26: Are all results required to be sent to SIEM or can results be presented in our own interface?

A26: Eligible technologies must provide a programmatically accessible output that can be forwarded to the SIEM, which will be Splunk. As stated in the Challenge.gov website.

Q27: Will we be able to apply updates for Policies and product on a weekly or bi-weekly basis?

A27: No, updates are not allowed after submitting the tool no later than June 30th, 2020.

Q28: Will Splunk be the SIEM solution?

A28: Yes, Splunk will be the SIEM for candidate tools to use.

Q29: How long will the evaluation duration be? If it is 7 months of activity, will we have the opportunity to do maintenance windows for patches, updates, etc.?

A29: As stated on the Challenge.gov website, licenses should be provided with an expiration date of 31 Dec 2020. As stated in Question #27 above, no patches/updates are allowed after the tool is submitted on or before June 30th, 2020.

Q30: Will network traffic generated from various network nodes (firewalls, Active directory, DNS) stored in a single unified SIEM (e.g Splunk) index log file or will there be multiple SIEM log files from different source network assets?

A30: Each technology will have a separate SIEM (Splunk) instance, and are required to log their alerts there. Outside logs will not be provided/available.

Q31: Is there an industry standard format being followed for how the network traffic fields are represented that are generated by these network nodes? OR should we NOT make any assumptions about the network field format presence?

A31: The latter—no assumptions about the network field format should be made.

Q32: In terms of clear text network traffic, what is the granularity of clear text network traffic. For example, will URIs and HTTP request/response headers be collected in the SIEM?

A32: This information is not given. The submitted network tools will have a mirrored copy of a SPAN and access to a SIEM, so a submitted participant technology could log events as desired.

Q33: Are we going to be given the certificates to decrypt the traffic or is the firewall inline where it can act as a proxy?

A33: Participant technologies will not be inline. The participant technologies will all be fed a mirrored copy of north-south (internal-external) & east-west (internal-internal) traffic from a single SPAN. Participants should expect a mix of encrypted and unencrypted traffic to be provided. The use of decryption technologies will not be permitted and are out of scope for the challenge.

Q34: If multiple configurations of Manufacturer A (MA) technology are submitted to the Prize Challenge…one as MAs official submission to the Prize Challenge and another through a partner organization, combining multiple manufacturer technologies as a full solution, is that allowed? Or will you only test each manufacturer's gear once?

For example, if #1: MA's official submission includes MA's hardware only as its solution and #2: MA was asked by a partner (i.e. systems integrator or value added reseller) to provide technology so they could submit a separate solution to include MA appliances and other manufacturer's appliances. Can MA submit #1 as their official Prize Challenge submission and also submit a note of consent, allowing MA's hardware to be used in #2? In the case of #2 winning the challenge, we acknowledge that as MA we could only be named the winner of #1 and not #2 under the rules of the Prize Challenge.

A34: A third party is allowed to submit a tool utilizing multiple manufacturers' appliances. However, the third party must have written consent from each manufacturer/product owner prior to submitting their solution. This consent must come in the form of an email that is submitted to the AI ATAC email address. If the third party's submission won, Manufacturer A would have no interest in the prize.

Q35: Which one of these best describes the solution:

a. An AI platform that correlates streams of security incidents into a composed and re-constructed high level attack campaign to support the SOC in attacker centric operations and security posture briefings
b. An IDS powered by AI that detects anomalies or malicious behavior in the network from a stream of network traffic, and produces a stream of alerts

A35: B is the correct answer. The detector does not have to rely solely on AI.

Q36: Should the output of the solution necessarily be a UDP stream of alerts (as is once mentioned in the whitepaper template)? Can it be a custom built usable and interactive user interface (for instance to showcase the campaigns progress evolution)?

A36: For the purposes of the competition, the output needed is a set of programmatically accessible alerts to be ingested by the SIEM. To quote the AI ATAC 2 call:

"Technologies will generate network alerts ingestible by the SIEM and using an alert format compatible with common SIEMs (e.g. Splunk)."

Splunk will be used for the SIEM.

If a user interface (UI) is part of the submission, the evaluation team may provide feedback to NAVWAR on the UI, but it is not a necessary component for and will not affect scoring for the purposes of the AI ATAC 2 competition results.

Q37: Is there a set environment for testing? Or can we assume a reasonable environment and deliver it in VMs, as well as providing specifications and scripts for its replication?

A37: There will be a setup environment for testing

You may deliver VMs with appropriate instructions as outlined in the call. The software and/or hardware components for the on-premises management server shall conform to the following options:

- Exported Virtual Machine Image (e.g., .ova, .qcow2, etc.) that can be run with a libvirt-compatible hypervisor (e.g. QEMU, XenServer, VMWare, Virtualbox, Emulab, etc.)
- Docker container package with comprehensive and clear setup documentation

- Hardware appliance

All submissions must be easily configurable to be on-premises only and must not make any connections to external cloud services

Include a license through 31 DEC 2020

Q38: Do we supply you with the SIEM as well? Or is there an existing SIEM that we need to integrate our solution with?

A38: Per the Judging Criteria of the Challenge posting, a SIEM will be provided. The SIEM will be a Splunk instance. Documentation must be included for integrating the submitted tool with Splunk. If the tool does not have native support for Splunk, please provide reasonable instructions for setting up the submitted tool with a Splunk forwarder.

Q39: In case of an existing SIEM, are there any:

a) **security applications running on it?**
b) **pre-defined detection rules / correlation queries installed?**
c) **APIs for accessing the SIEM? If so, is it documented?**

A39: Answers are as follows:

a) No
b) No
c) Splunk will be the SIEM

Q40: Can we use a commercial signature based detection system (e.g. Suricata) as an augmentation to our AI/ML solution, making it a bundle of security technologies?

A40: Open source technologies used as part of the submission technology are permitted and should be documented in the white paper portion of the submission. Partnerships with proprietary technologies is permitted, but this requires written consent by the third parties for the use of these capabilities, along with sufficient permissions to allow the Government to evaluate the submission for the Prize Competition.

Q41: What data sources can we assume to have in the environment? In the challenge statement, only a passive network tap (i.e. raw PCAPs) have been mentioned. Are the following absent in the environment? If so, can we deliver these along with the solution?

a) **Network Data (e.g. Netflow, NIDS, FW, DLP etc.)**
b) **Auth logs (e.g. Active Directory, Single Sign On, etc.)**
c) **Host Data (e.g. EDR, EPP, HIDS, Syslogs, Host Audit Data etc.)**
d) **Web Data (e.g. Web Proxy, Web filter, etc.)**

A41: Answers are as follows:

a) Network-level alerting components (NIDS, FW, DLP) will not be present but are the subject of the test (what should be submitted).

   If a submission requires a form of netflow data, the netflow sensor should be a part of the submission and should function as desired with the other (network detection) components of the submission.

   More generally, network level detection technologies can partner with third-party technologies (e.g. load balancers), but this requires written consent by the third parties for the use of these complementary capabilities. The network detection technology is what is eligible for award and evaluation. The third party is not eligible for award in this case. These third parties could be used by multiple eligible network detection technologies in the AI ATAC 2 Challenge (with their consent for each).

b) To quote the call:

   “Network services include (minimally) a Firewall, Active Directory, Domain Name Service, and Dynamic Addressing (DHCP), and the network data will contain traffic from these services.”

c) To quote the call:

   “ The candidate product will not receive information directly from or rely upon the host endpoint security service.”

   Presence or absence of host logging (e.g., syslogs, audit data) collection will not be disclosed.

d) To quote the call:

"All north/south and east/west traffic is observable with a mixture of encrypted and unencrypted traffic."

Q42: Is the end of June a hard deadline?

A42: Yes, per the Challenge guidelines all submissions are due NLT June 30th, 2020 at 5PM Eastern

Q43: Is the current situation with the pandemic introducing any changes to the challenge funding or deadline?

A43: Due to COVID-19 the challenge submission date was already changed from May 29th, 2020 to June 30th, 2020 to allow participants additional time to submit their entry. There is no change in the prize amount.

Q44: Can the vendor send their server pre-configured with software or does it have to be bare metal and install happens on site?

A44: Both are appropriate submission configurations provided adequate instructions are included to stand-up the technology in the research range and ensure proper functionality.

Q45: The license(s) to operate the technology on a 10Gb/s bandwidth network through 31 DEC 2020 on multiple VMs simultaneously: Does this mean multiple independent instances running?

A45: We plan to use a single instance in the evaluation, but licenses for a multiple (~5) instances is preferred, in particular so multiple engineers can work on proper configuration and debugging in parallel.

Q46: 1,500 node network: Network infrastructure devices or end user devices?

A46: "Roughly 1,500 nodes" refers to the total number of infrastructure and endpoint devices.